

AD-A260 279



213375-3-T

DTIC
ELECTE
FEB 16 1993
S C D

2

FOCAL-PLANE ALIGNMENT SENSING

R. G. PAXMAN
J. H. SELDIN

Environmental Research Institute of Michigan
P.O. Box 134001
Ann Arbor, MI 48113-4001

FEBRUARY 1993

Final Report for Period 3 April 1992-30 September 1992

Distribution: Approved for Public Release; Distribution is Unlimited.

Phillips Laboratory
Kirtland Air Force Base, NM 87117-6008

Contract No. DLA900-88-D-0392

93-02835



P.O. Box 134001
Ann Arbor, MI 48113-4001



93 2 12 151

PREFACE

This report documents the results of the first phase of a proposed two-phase study called "Focal-Plane Alignment Sensing" performed by the Optical & Infrared Science Laboratory of the Environmental Research Institute of Michigan, Ann Arbor, Michigan, for Phillips Laboratory, Kirtland Air Force Base, New Mexico, during the period 03 April 1990 to 30 September 1992. This study was performed under Delivery Order 0028 within the Infrared Information Analysis Center (IRIA) program, contract number DLA900-88-D-0392, for which the Defense Electronic Supply Center (DESC), Dayton, Ohio, serves as the contracting agency. The Phillips Laboratory project managers were Nancy Miller and Christopher DeHainaut. The Principal Investigator was Richard G. Paxman. The authors of this report are Richard G. Paxman and John H. Seldin.

DTIC QUALITY INSPECTED 3

Accession For	
NTIS SECRET	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
DTIC	Special
A-1	

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE February 1993		3. REPORT TYPE AND DATES COVERED Final - 3 April 1992 to 30 Sept. 1992
4. TITLE AND SUBTITLE Focal-Plane Alignment Sensing			5. FUNDING NUMBERS DLA900-88-D-0392	
6. AUTHOR(S) Richard G. Paxman and John H. Seldin				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Environmental Research Institute of Michigan P.O. Box 134001, Ann Arbor, MI 48113-4001			8. PERFORMING ORGANIZATION REPORT NUMBER 213375-3-T	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Phillips Laboratory Kirtland Air Force Base, NM 87117-6008			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The method of phase diversity has been investigated for use in sensing piston and tilt misalignments in the Multipurpose Multiple Telescope Testbed (MMTT), a phased-array telescope in operation at Phillips Laboratory, Kirtland Air Force Base. Misalignment parameters are estimated indirectly from focal-plane imagery. Accordingly, the hardware requirements for this approach are modest. Furthermore, the method relies on an external reference (the object being imaged) so that it is less susceptible to systematic errors than are direct-sensing approaches. Regularization is used to combat noise sensitivity, and strategies for selecting a regularization parameter have been developed. The probability of convergence to within a prescribed RMS phase error (operational convergence) was quantified for noisy imagery using Monte Carlo experiments and was found to be acceptable. Simulations were also performed using three diversity channels, in which the fidelity of estimates was shown to improve and the probability of operational convergence was observed to dramatically increase.				
14. SUBJECT TERMS Phase diversity Regularization Multipurpose Multiple Telescope Testbed (MMTT)			15. NUMBER OF PAGES 50	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified
20. LIMITATION OF ABSTRACT Unlimited				

Contents

Preface	iii
List of Figures	vi
1.0 Introduction	1
1.1 Statement of Phase-Diversity Problem	1
1.2 Summary of Research and Conclusions	3
2.0 Regularization in Phase Diversity	9
2.1 Regularized Objective Function	9
2.1.1 Constrained Least Squares Problem	9
2.1.2 Derivation of Regularized Objective Function	11
2.2 Selection of Regularization Parameter	12
2.2.1 Existing Strategies in Image Deblurring	12
2.2.2 CLS Parameter Selection	14
2.2.3 Wiener Parameter Selection	16
3.0 MMTT Operational Convergence	19
3.1 Review of Monte-Carlo Approach	19
3.2 Experiment Design	20
3.3 Experimental Assumptions	23
3.4 Review of Noiseless Experiment	24
3.5 Operational Convergence with Noise	27
3.6 Experimental Results	28
3.7 Monte-Carlo Experiment Summary	31
4.0 Multiple Diversity Frames	33
4.1 Potential Advantages of Multiple Diversity Frames	33
4.2 Three-Frame Diversity Experiment	34
4.3 Three-Frame Probability of Operational Convergence	35
4.4 Summary of Multiple-Frame Phase Diversity	40
Appendix A. Applied Phase Diversity Software Upgrades	A-1

List of Figures

3-1	Estimate of PDF for MMTT RMS Misalignment Error	22
3-2	MMTT Probability of Operational Convergence with No Noise	26
3-3	Probability of Operational Convergence Versus Convergence Threshold	30
3-4	MMTT Probability of Operational Convergence	31
4-1	Three-Frame Phase Diversity	37
4-2	MMTT Two- and Three-Frame Probability of Operational Convergence	39

1.0 INTRODUCTION

The Multipurpose Multiple Telescope Testbed (MMTT), designed and built under the direction of researchers at Phillips Laboratory, Kirtland Air Force Base, is the first wide field-of-view phased-array imaging telescope to be built. An elaborate sensor system, consisting of local laser interferometers and position detectors, is used to estimate piston and tilt misalignments. This misalignment-sensing system works but has two potential problems. First, the sensor optics are complex because many individual measurements are required to determine piston and tilt misalignments among all four subtelescopes. Second, the sensor is subject to bias errors arising from drift and misalignments in the sensor optics.

We have investigated a technique known as phase diversity as a complement or an alternative to the existing direct aberration-sensing system. Phase diversity requires the collection of (at least) two incoherent images: one in the focal plane and a second in a location that intentionally introduces a known amount of defocus. The additional hardware required for the technique is modest and includes a simple beam splitter and a second detector array. No local sources, such as lasers, are required. In addition, the method relies on an external reference – the object being imaged. Therefore the method should be less susceptible to bias errors than direct-sensing methods.

Our research into the use of phase diversity in the MMTT system includes a two-year effort documented in a report entitled *Applied Phase Diversity* [1.1]. The present report documents a six-month effort that represents the first phase of a proposed two-phase follow-on effort entitled *Focal-Plane Alignment Sensing*. Therefore, the results contained in this report build heavily on the previous two-year effort. Although we continue to conclude that phase diversity provides a viable alignment-sensing approach, the MMTT program at Phillips was fundamentally redirected during the course of this study and it is unlikely that the second phase of the current effort will be undertaken.

1.1 STATEMENT OF PHASE-DIVERSITY PROBLEM

The concept of phase-diversity measurements in optical systems was first introduced by Gonsalves [1.2, 1.3], who suggested using phase diversity to estimate atmospheric aberrations in earth-based telescopes. An incoherently illuminated object is imaged

with an aberrated imaging system. In general, these aberrations could derive from misalignments, improper mirror figure, off-axis aberrations, or atmospheric turbulence. The conventional image will then suffer from these aberrations. A second image that is intentionally defocused by a known amount is easily collected with a beam splitter and a second detector array. This “diversity” image will be degraded by both the system aberrations and the quadratic phase aberration associated with the defocus. The task is to jointly estimate both the object and the system aberrations from the two images. This task was recently formalized using estimation theory [1.4]. Since the seminal work by Gonsalves, others have studied the use of phase diversity in phased-array systems [1.1,1.5-1.7] and for imaging through the atmosphere [1.8-1.12].

So that this document will be self-contained, we now review the formal statement of the phase-diversity problem. This review also serves to introduce mathematical notation that is used throughout this report. Consider the space-invariant (isoplanatic) incoherent imaging equation:

$$d_k(x) = f * s_k(x) + n_k(x) \quad (1-1)$$

$$= g_k(x) + n_k(x), \quad k = 1, \dots, K \quad (1-2)$$

where f is the object, s_k is the k th *point-spread function* (PSF), g_k is the ideal (noiseless) image, n_k is the detector noise, d_k is the detected image, and the asterisk denotes convolution. The subscript k indexes the individual data sets. Traditionally, there has been two data sets, a conventional image and a diversity image. In general, however, $K \geq 2$. For incoherent imaging, the PSF is the modulus squared of the coherent impulse response function, which in turn is just the Fourier transform of the coherent transfer function:

$$s_k(x) = |\mathcal{F}^{-1}\{|H_k(u)|e^{i\phi(u)}e^{i\theta_k(u)}\}|^2, \quad (1-3)$$

where $|H_k|$ is the binary pupil function, ϕ is the unknown aberration function, θ_k is the k th phase diversity, and the operator $\mathcal{F}^{-1}\{\cdot\}$ takes the inverse Fourier transform of the argument. It is often convenient to parameterize the phase aberration function

$$\phi(u) = \sum_{j=1}^J \alpha_j \phi_j(u), \quad (1-4)$$

where $\{\alpha_j\}$ is the set of aberration parameters and $\{\phi_j\}$ is a convenient basis set. In the phased-array misalignment problem, for example, it is natural to use piston and tilt parameters. The set of parameters can be ordered in a single parameter vector, α . We can now state the problem. Given the set of detected images, $\{d_k\}$, the set of phase-diversity functions, $\{\theta_k\}$, and the set of binary pupil functions, $\{|H_k|\}$, jointly estimate the object, f , and the phase aberration, ϕ , (or the parameter vector α).

1.2 SUMMARY OF RESEARCH AND CONCLUSIONS

The goal of the research was to investigate the use of the method of phase diversity for the problem of sensing MMTT misalignments. Research activities included developing regularization techniques to combat noise amplification, quantifying the convergence of the algorithms, investigating performance when additional diversity images are collected ($K > 2$), and upgrading and refining the Applied Phase Diversity (APD) software package.

We have observed that noise amplification poses a significant problem both in the nonlinear optimization procedure and in the final estimates. In the previous effort [1.1] a regularized objective function was derived to combat this problem. In Section 2.1 a refined version of this derivation is presented along with additional insight into the method. The problem of regularization parameter selection is discussed in Section 2.2. Strategies of regularization parameter selection in conventional deblurring problems are surveyed and provide guidance for our more demanding problem of *joint* estimation of object and aberrations. A generalization of the Constrained Least Squares (CLS) strategy yields a well-defined prescription for adaptively selecting the regularization parameter, based upon knowledge of the detector noise level. This strategy is adaptive in the sense that the regularization parameter is a function of location in aberration parameter space. In practice we expect that this variation will be quite mild. A second strategy for regularization parameter selection derives from a restatement of the problem in which the scene is no longer deterministic, but is treated as a stationary random process. A maximum *a posteriori* formulation of the problem will then give an objective function with the same form as in the deterministic case, but with the parameter now being the ratio of the power spectra of the noise to that of the scene. This is analogous to the parameter found in Wiener filtering.

Whereas the class of problems for which there will be *a priori* knowledge about the power spectra of both the noise and the object scene is significantly restricted, this analysis provides insight into the role of the regularization parameter, gives a sanity check on the order of magnitude of the parameter, and will be appropriate for some classes of objects, such as homogeneous terrains.

In the previous effort we defined *operational convergence* to be achieved when the optimization sequence returns an aberration estimate that is within a given tolerance of the true solution, in a RMS sense. This definition eliminates the need to identify local or global maxima, since operationally we only care about the proximity of the estimate to the true solution. In Section 3 we review an elaborate Monte-Carlo experiment that quantifies the probability of operational convergence for the MMTT in the absence of noise. In this case the probability of operational convergence is found to be approximately 94%. Under the current effort these experiments were extended to the case in which noise was present. For noise having a standard deviation of 0.5% of the peak pixel in the image, the probability of operational convergence was found to be a little more than 50%. This level of noise is somewhat greater (or more pessimistic) than could be achieved if the full dynamic range of a good CCD were used. The probability of operational convergence as a function of RMS misalignment is plotted in Figure 3-4 (page 31). Whereas we have shown that the likelihood of algorithm operational convergence for the worst-case scenario of no *a priori* knowledge of misalignment is smaller in the presence of noise, we feel that the probabilities for startup and tracking is sufficiently high to make this an encouraging result.

In Section 4 we present simulation results in which the number of diversity frames has been increased ($K = 3$). Simulated imagery (Figure 4-1, page 36) provides evidence that under very noisy conditions an additional diversity image can significantly improve estimates for both the object and aberrations. Moreover, the probability of operational convergence is seen to improve dramatically with an additional diversity image. The 3-frame operational convergence in the presence of noise was found to be over 90%, as shown in Figure 4-2 (page 37).

Appendix A presents a variety of upgrades to the APD software that were accomplished in this effort. These include the implementation of a closed-form expression for the gradient of the objective function, the ability to operate on multiple diversity frames, the option to output a restored object image when regularization is used, and

certain improvements to insure numerical accuracy.

The milestones and results achieved in this and the previous [1.1] research efforts suggest that the prospect of using phase diversity to sense misalignments in phased-array telescopes is good. A variety of computational tools have been developed to reduce the computational expense and accommodate real data from scenes. There is strong evidence that the method is robust to noise. Moreover, the probability of operational convergence is sufficiently high to suggest the use of phase diversity in start-up or tracking modes.

REFERENCES

- [1.1] R.G. Paxman, J.H. Seldin, and P.P. Sanchez, "Applied Phase Diversity," final technical report to Phillips Laboratory (Kirtland Air Force Base), ERIM Report No. 213390-3-F, March 1992.
- [1.2] R.A. Gonsalves and R. Childlaw, "Wavefront sensing by phase retrieval," in *Applications of Digital Image Processing III*, A.G. Tescher, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **207**, 32-39 (1979).
- [1.3] R.A. Gonsalves, "Phase retrieval and diversity in adaptive optics," Opt. Eng. **21**, 829-832 (1982).
- [1.4] R.G. Paxman, T.J. Schulz, and J.R. Fienup, "Joint estimation of object and aberrations using phase diversity," J. Opt. Soc. Am. A **9**, 1072-1085 (1992).
- [1.5] R.G. Paxman and J.R. Fienup, "Optical misalignment sensing and image reconstruction using phase diversity," J. Opt. Soc. Am. A **5**, 914-923 (1988).
- [1.6] R.G. Paxman and S.L. Crippen, "Aberration correction for phased-array telescopes using phase diversity," in *Digital Image Synthesis and Inverse Optics*, Proc. Soc. Photo-Opt. Instrum. Eng. **1351**, 787-797 (1990).
- [1.7] N. Miller, "Nonlinear Optimization and Neural Networks for Phase Diversity: Simulations and Experiment," Ph.D. thesis (University of New Mexico, Albuquerque, NM, 1992).
- [1.8] J.A. Högbom, "On the intensity distribution over the focal volume," in *High Spatial Resolution Solar Observations*, Proceedings of the 10th Sacramento Peak Summer Workshop, Sunspot, New Mexico (1988).
- [1.9] J.A. Högbom, "Reconstruction from focal volume information," in *Solar and Stellar Granulation*, R.J. Rutten and G. Severino, eds. (Kluwer Academic Publishers, Hingham, MA, 1989).

- [1.10] S.R. Restaino, "Wavefront sensing and image deconvolution of solar data," *Applied Optics* **35**, 7442-7449 (1992).
- [1.11] R.G. Paxman, T.J. Schulz, and J.R. Fienup, "Phase-diverse speckle interferometry," in *Signal Recovery and Synthesis IV, 1992 Technical Digest Series 11* (Optical Society of America, Washington DC, 1992).
- [1.12] R.G. Paxman and J.H. Seldin, "Fine-resolution imaging of solar features using phase-diverse speckle imaging," in *Real Time and Post-Facto Solar Image Correction*, Proceedings of the 13th Sacramento Peak Summer Workshop, Sunspot, New Mexico (1992).

2.0 REGULARIZATION IN PHASE DIVERSITY

In the case of additive Gaussian noise, the derivation of the phase-diversity objective function involves substituting an intermediate estimate for the object (given the aberrations) into the log-likelihood function for the joint estimation of the object and aberrations. The intermediate estimate has the form of an inverse filter estimate, which is known to be ill-conditioned, or intolerably sensitive to noise. We have observed that the objective function develops a rough, fractal-like texture when noise is introduced into simulations [2.1]. Of course, such a texture will pose severe problems for almost any nonlinear optimization algorithm. We believe that this fractal-like texture is the direct result of noise amplification induced by the inverse filter.

The problem of noise amplification that arises in conventional image deblurring problems has often been addressed by means of *regularization*. In the broad sense of the word, regularization is a procedure to bridle noise amplification by limiting the class of feasible solutions. The restriction of the class of feasible solutions represents some type of *a priori* knowledge that is brought to bear on the problem, explicitly or implicitly. In a previous research effort, a regularized objective function was derived for the phase-diversity problem. In the research effort described by this report, the regularized objective function was refined and strategies for the selection of a regularization parameter were developed. In the next section we highlight the derivation of the regularized objective function. Section 2.2 presents two strategies for the selection of a regularization parameter from the data.

2.1 REGULARIZED OBJECTIVE FUNCTION

In this section we highlight the derivation of the regularized objective function, which is carried out in detail in [2.2]. We also provide new insight regarding the basis of the theory. There is a considerable body of literature on the subject of regularization for the problem of image deblurring. This literature serves as a guide in the generalization to our problem, in which both the object and the aberrations are jointly estimated.

2.1.1 Constrained Least Squares Problem

As with the unregularized estimate [2.2,2.3], we wish to maximize the log-likelihood function, L . However, we now introduce a constraint on the object estimate that is

intended to bridle the noise amplification. Formally,

$$\max_{f, \alpha} \{L = -\frac{1}{N^2} \sum_{k=1}^K \|D_k(u) - F(u)S_k(u)\|^2\} \quad (2-1)$$

$$\text{subject to} \quad \|Q(u)F(u)\|^2 = C_1, \quad (2-2)$$

where $\|\cdot\|^2$ represents the norm squared (sum over u of squared magnitudes), C_1 is a constant, and $Q(u)$ is a multiplicative regularizing operator that influences the role of the constraint. A common choice of $Q(u)$ is unity. In this case the constraint simply fixes the norm of the object estimate, thus preventing wild oscillations that otherwise might occur due to noise amplification. The function $Q(u)$ can also be selected to enforce a smoothness constraint on the object estimate [2.4]. This formulation of the problem is analogous to a deblurring method proposed by Hunt [2.4] and may be properly referred to as a constrained least-squares method. The prior knowledge being invoked is knowledge that the weighted norm of $F(u)$ is fixed and equal to C_1 . It is difficult to imagine many realistic problems in which such knowledge will be available.

Consider an alternative problem statement:

$$\min_{f, \alpha} \{\|Q(u)F(u)\|^2\} \quad (2-3)$$

$$\text{subject to} \quad \sum_{k=1}^K \|D_k(u) - F(u)S_k(u)\|^2 = C_2, \quad (2-4)$$

where C_2 is a known constant. In this formulation, the constraint restricts the sum of squared differences between the noiseless image and the detected data to be a constant. For the case of additive Gaussian noise, this expression is just the detector noise and we are suggesting that this value (C_2) is fixed and known. It is quite reasonable to expect that one might know the level of detector noise *a priori*. The goal here is to find the object with the minimum weighted norm, subject to a fixed amount of detector noise. Philosophically, this viewpoint is quite different from the

constrained maximum-likelihood problem given in Eqs. (2-1) and (2-2). However, both problem statements can be analyzed using the method of Lagrange multipliers and yield the same auxiliary function:

$$A(f, \alpha) \equiv \sum_{k=1}^K \|D_k(u) - F(u)S_k(u)\|^2 + \lambda \|Q(u)F(u)\|^2. \quad (2-5)$$

Therefore, these two problem statements will give the same form for the objective function. Whereas the constrained maximum-likelihood formulation may be more satisfying philosophically, the minimum-norm formulation utilizes *a priori* knowledge that is more accessible. One can view the minimum-norm formulation as a means of converting *a priori* information about detector noise (C_2) into *a priori* information about the weighted object norm (C_1). The Lagrange multiplier λ must be adjusted to satisfy the constraint being used, so that there is a coupling between the values C_1 , C_2 , and λ .

It is worth noting that the auxiliary function given in Eq. (2-5) is often the point of departure for regularization schemes. Tikhonov [2.5] and others have viewed the second term in Eq. (2-5) as a regularizing term with λ serving as a regularizing parameter that can be continuously varied to adjust the degree of smoothing in the estimate. However, Tikhonov regularization applies to the case for which only f is to be estimated and α , or equivalently S , is considered known.

2.1.2 Derivation of Regularized Objective Function

The objective function is formed by first temporarily assuming that the aberration parameters are fixed and finding an expression for the regularized object. This is accomplished by taking the derivative of A with respect to the real and imaginary parts of F and setting to zero. The regularized object estimate is expressed by

$$F_R(u) = \frac{\sum_{k=1}^K D_k(u)S_k^*(u)}{\lambda|Q(u)|^2 + \sum_{l=1}^K |S_l(u)|^2}. \quad (2-6)$$

Substitution of the regularized object estimate into the auxiliary function yields the regularized objective function:

$$L_R(f, \alpha) = \sum_{u \in X} \frac{|\sum_{j=1}^K D_j S_j^*|^2}{\lambda |Q|^2 + \sum_{l=1}^K |S_l|^2}, \quad (2.7)$$

where an inconsequential constant term has been dropped and the u -dependence has been suppressed. The details of the derivations of Eqs. (2-6) and (2-7) are found in [2.2]. We have observed that the regularized objective function no longer has a fractal-like texture and accommodates nonlinear optimization. Furthermore, a closed-form expression for the gradient of the regularized objective function has been derived [2.2].

2.2 SELECTION OF REGULARIZATION PARAMETER

Having developed a regularized objective function, we now turn to the important question of selection of the regularization parameter. Recall that increasing the regularization parameter reduces noise amplification, but only at the expense of spatial resolution. We need some strategy for selecting the regularization parameter that appropriately balances these two effects.

2.2.1 Existing Strategies in Image Deblurring

Researchers interested in applying regularization concepts to conventional image deblurring problems have developed a variety of strategies for the selection of the regularization parameter value. Regularization parameter selection continues to be a subject of active research in this area. We briefly enumerate some of the more prominent strategies that have been proposed, as they provide a guide for us in developing strategies for our problem. The first approach is referred to variously as Constrained Least Squares (CLS) [2.4,2.6,2.7], Chi-square [2.8], and the discrepancy principle [2.9]. The CLS method is widely known and has been popular throughout the use of regularization techniques. Similar in spirit to our constrained least squares problem formulation, this method assumes that *a priori* information regarding the level of detector noise is available. A second strategy for regularization parameter selection, due to Wahba, is referred to as Equivalent Degrees of Freedom (EDF) [2.7,2.8,2.10]. This method is slightly more sophisticated than the CLS method, however it still

relies upon knowledge of detector noise. A third approach seeks to estimate the regularization parameter from self-consistency criteria using only the detected data. No *a priori* information about the noise level is presupposed. This approach can be derived from the *leave one out* principle. This class of strategies includes ordinary and generalized cross validation (CV) [2.7,2.8,2.11]. Cross-validation methods are computationally burdensome but because no *a priori* information about the noise level is required, they are currently a subject of intense research. There have even been efforts to estimate both the regularization parameter and the regularizing operator from the data alone [2.12]. A fourth strategy often goes unmentioned but has been used very often historically. This is the interactive choice of the parameter. In this method, the image interpreter interactively adjusts the parameter to the level that gives the most pleasing or useful image estimate. This method accommodates the reservoir of informal *a priori* information that resides within the interpreter. Whereas this method is suspect when the image content is unfamiliar to the interpreter, it has significant merit for cases in which image components can be recognized. A final strategy for the selection of the regularization parameter actually requires restating the problem so as to treat the object as a realization from a stationary random process. The maximum *a posteriori* (MAP) solution to this problem is well known as a Wiener estimate [2.13] and the regularization parameter can be identified as the ratio of noise variance to the object variance [2.14]. This approach demands even more *a priori* information since object statistics must be known in addition to noise statistics. For certain homogeneous terrains, such information may be available.

Each of the strategies surveyed could probably be generalized to accommodate the problem of *joint* estimation of object and aberrations. Rather than generalize each method and then determine the relative performance of each, we choose to concentrate on two of the more straightforward methods. The two that we choose to generalize are the CLS and the Wiener methods. These two methods demonstrate that noise amplification can be suppressed in a reasonable fashion for the joint estimation problem. The remaining methods could be investigated given additional resources.

2.2.2 CLS Parameter Selection

In many circumstances, the level of detector noise will be available *a priori*. When this is true, then the CLS strategy for selecting a regularization parameter in the case of deblurring problems can be generalized to the problem of joint estimation of object and aberrations. We assume that the noise is independent and identically distributed at each detector element with a variance of σ_n^2 . It is straightforward to identify the constant C_2 defined in Eq. (2-4).

$$C_2 = \sum_{k=1}^K ||D_k(u) - F(u)S_k(u)||^2 \quad (2-8)$$

$$= KN^2\sigma_n^2, \quad (2-9)$$

where N^2 is the number of detector elements in a single detector array. By substituting the regularized object estimate, F_R (Eq. (2-6)), into Eq. (2-8), we get an expression that *implicitly* defines λ for each aberration vector estimate.

$$C_2 = \sum_{u \in \mathcal{X}} \sum_{k=1}^K |D_k(u) - F_R(u)S_k(u)|^2 \quad (2-10)$$

$$= \sum_{u \in \mathcal{X}} \sum_k \left| D_k - S_k \frac{\sum_j D_j S_j^*}{\lambda|Q|^2 + \sum_l |S_l|^2} \right|^2 \quad (2-11)$$

$$= \sum_{u \in \mathcal{X}} \sum_k \left| \frac{D_k(\lambda|Q|^2 + \sum_l |S_l|^2) - S_k \sum_j D_j S_j^*}{\lambda|Q|^2 + \sum_l |S_l|^2} \right|^2 \quad (2-12)$$

$$= \sum_{u \in \mathcal{X}} \sum_k \frac{|D_k(\lambda|Q|^2 + \sum_l |S_l|^2)|^2 + |S_k \sum_j D_j S_j^*|^2 - (D_k^*(\lambda|Q|^2 + \sum_l |S_l|^2) S_k \sum_j D_j S_j^* + C.C.)}{(\lambda|Q|^2 + \sum_l |S_l|^2)^2} \quad (2-13)$$

where the u -dependence has been suppressed, the summations over j , k , and l run from 1 to K , and $C.C.$ represents a term that is the complex conjugate of the preceding term. Notice that the summation over k only affects the numerator in Eq. (2-13). We now distribute the k -summation over each term in the numerator to get

$$\begin{aligned}
\text{numerator} &= \sum_k |D_k|^2 (\lambda |Q|^2 + \sum_l |S_l|^2)^2 + \sum_k |S_k|^2 \sum_j D_j S_j^* |^2 \\
&\quad - \sum_k D_k^* S_k (\lambda |Q|^2 + \sum_l |S_l|^2) \sum_j D_j S_j^* - \sum_k D_k S_k^* (\lambda |Q|^2 + \sum_l |S_l|^2) \sum_j D_j^* S_j \\
&\hspace{25em} (2-14)
\end{aligned}$$

$$\begin{aligned}
&= \sum_k |D_k|^2 (\lambda |Q|^2 + \sum_l |S_l|^2)^2 - \lambda |Q|^2 \sum_j D_j S_j^* |^2 - (\lambda |Q|^2 + \sum_l |S_l|^2) \sum_j D_j S_j^* |^2 \\
&\hspace{25em} (2-15)
\end{aligned}$$

Reintroducing the denominator and factoring out common factors where appropriate, we get

$$C_2 = \sum_{u \in \mathcal{X}} \sum_k |D_k|^2 - \sum_{u \in \mathcal{X}} \frac{|\sum_j D_j S_j^*|^2}{\lambda |Q|^2 + \sum_l |S_l|^2} - \sum_{u \in \mathcal{X}} \frac{\lambda |Q|^2 |\sum_j D_j S_j^*|^2}{(\lambda |Q|^2 + \sum_l |S_l|^2)^2} \quad (2-16)$$

Equation (2-16) implicitly defines the parameter λ for specific OTF estimates, $\{S_k\}$, or equivalently for the aberration vector estimate, α . Thus the value of λ that appears in the regularized objective function should properly be regarded as a function of the aberration parameter vector:

$$L_R(f, \alpha) = \sum_{u \in \mathcal{X}} \frac{|\sum_{j=1}^K D_j S_j^*|^2}{\lambda(\alpha) |Q|^2 + \sum_{l=1}^K |S_l|^2} \quad (2-17)$$

In order to strictly evaluate the constrained/regularized objective function for a specific aberration vector, one must implicitly solve for $\lambda(\alpha)$ by applying an iterative method (such as Newton's method) to Eq. (2-16) and then use this solution to evaluate Eq. (2-17). Whereas in the conventional deblurring problem λ is fixed, in the joint-estimation problem λ *adapts* as the aberration vector varies. In practice, λ is expected to be a slowly varying function, so that its value may only need to be updated occasionally during an optimization sequence.

We now have a well-defined prescription for adaptively selecting the regularization parameter, based on knowledge of the detector noise level, in the joint estimation of the object and aberrations.

2.2.3 Wiener Parameter Selection

To this point we have modeled the object as a deterministic function. If we view the object as a single realization of a random process, then the problem statement must be modified. We now assume that the scene is a stationary, Gaussian stochastic process with a known power spectrum, $P(u)$. A maximum *a posteriori* (MAP) estimator [2.15] for the joint estimation of object and aberrations can then be constructed. To do so we use a uniform prior probability for the aberration parameters (since we treat these as deterministic). An intermediate assumption that the aberrations are fixed yields the MAP estimate for the object, also known as the Wiener estimate [2.13]:

$$F_W = \frac{\sum_k D_k S_k^*}{KN^2 \sigma_n^2 / P + \sum_l |S_l|^2} \quad (2-18)$$

Equation (2-18) is identical to the regularized object estimate, expressed in Eq. (2-6), if the ratio of power spectra is substituted for the product of the regularization parameter and the modulus squared of the regularizing operator,

$$\lambda |Q(u)|^2 \rightarrow KN^2 \sigma_n^2 / P(u) \quad (2-19)$$

This correspondence has been noted by numerous authors, including [2.14]. Substitution of Eq. (2-18) into the expression for the *a posteriori* probability gives an objective function that has no explicit dependence upon the object:

$$L_W(f, \alpha) = \sum_{u \in \mathcal{X}} \frac{|\sum_{j=1}^K D_j S_j^*|^2}{KN^2 \sigma_n^2 / P + \sum_{l=1}^K |S_l|^2} \quad (2-20)$$

We see that the objective functions expressed in Eq. (2-17) and Eq. (2-20) are also identical, given the correspondence prescribed in Eq. (2-19). Thus the MAP or Wiener perspective has given a new strategy for estimating both the regularization parameter and the regularization operator and has provided an additional interpretation for these entities. It is worth noting that, whereas the CLS prescription yields a regularization parameter that adapts as the aberration parameters are varied, the Wiener viewpoint prescribes a regularization parameter that is constant throughout the optimization sequence. This suggests that the CLS regularization parameter adapts very slowly, as expected.

REFERENCES

- [2.1] R.G. Paxman and J.R. Fienup, "Optical misalignment sensing and image reconstruction using phase diversity," J. Opt. Soc. Am. A **5**, 914-923 (1988).
- [2.2] R.G. Paxman, J.H. Seldin, and P.P. Sanchez, "Applied Phase Diversity," final technical report to Phillips Laboratory (Kirtland Air Force Base), ERIM Report No. 213390-3-F, March 1992.
- [2.3] R.G. Paxman, T.J. Schulz, and J.R. Fienup, "Joint estimation of object and aberrations by using phase diversity," J. Opt. Soc. Am. A **9**, 1072-1085 (1992).
- [2.4] B.R. Hunt, "The application of constrained least squares estimation to image restoration by digital computer," IEEE Trans. Comput. **C-22**, 805-812 (1973).
- [2.5] A.N. Tikhonov and V.Y. Arsenin, *Solutions of Ill-Posed Problems*, Winston, Minneapolis (1977).
- [2.6] A.K. Katsaggelos, "Iterative image restoration algorithms," Opt. Eng. **28**, 735-748 (1989).
- [2.7] N.P. Galatsanos and A.K. Katsaggelos, "Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation," IEEE Trans. on Image Processing **1**, 322-336 (1992).
- [2.8] A.M. Thompson, J.C. Brown, J.W. Kay, and D.M. Titterington, "A study of methods of choosing the smoothing parameter in image restoration by regularization," IEEE Trans. Pattern Anal. Machine Intell. **13**, 326-338 (1991).
- [2.9] C.K. Rushforth, "Signal restoration and Fredholm integral equations," in *Image Recovery: Theory and Applications*, Henry Stark, ed., (Academic Press, Orlando, Florida, 1987).
- [2.10] G. Wahba, "Bayesian 'confidence intervals' for the cross-validated smoothing splines," J. Roy. Stat. Soc. B **45**, 651-667 (1977).

- [2.11] G.H. Golub, M.Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometr.* **21**, 215-223 (1979).
- [2.12] S.J. Reeves and R.M. Mersereau, "Optimal estimation of the regularization parameter and stabilizing functional for regularized image restoration," *Opt. Eng.* **29**, 446-454 (1990).
- [2.13] C.W. Helstrom, "Image restoration by the method of least squares," *J. Opt. Soc. Am.* **57**, 297-303 (1967).
- [2.14] N.B. Karayiannis and A.N. Venetsanopoulos, "Regularization theory, MAP estimation and Wiener filtering in image restoration: points of tangency and areas of intersection," in *Proc. Int. Symposium on Circuits and Systems*, IEEE, 835-838 (1987).
- [2.15] P.L. Van Trees, *Detection, Estimation, and Modulation Theory - Part I*, (Wiley, New York, 1968).

3.0 MMTT OPERATIONAL CONVERGENCE

In a previous MMTT phase-diversity research program [3.1], a Monte-Carlo experiment was used to estimate the probability of operational convergence for phase-diversity and for expected MMTT misalignments. Estimates were given for the probability of estimating a misalignment with a certain Root-Mean Square (RMS) phase error to within a prescribed tolerance, our definition of *operational convergence*. These experiments were performed for the case of no noise, and we discuss here the same experiments performed in the presence of additive Gaussian detector noise, using the Wiener regularization strategy discussed in the previous section. In this section, the Monte-Carlo approach to estimating operational convergence is reviewed, and a summary of the original, noiseless experiment is given. The assumptions made in determining the appropriate level of additive noise are also discussed and the results of the probability of operational convergence experiment in the case of noise are presented.

3.1 REVIEW OF MONTE-CARLO APPROACH

An undesirable attribute of a nonlinear optimization algorithm, such as is required in the phase diversity problem, is the potential for stagnation at a sub-optimal solution in the search for the true solution. Ideally, one would like to comprehend the surface contour of the phase-diversity objective function to infer the concentration of local maxima and the volume in parameter space for which convergence to the global maximum is guaranteed. Moreover, one would like to know this contour as a function of the true parameter values. For the MMTT, this amounts to sampling the objective function in an 18-dimensional parameter space (9 true parameters and 9 estimated parameters). Such a task is obviously intractable. We took a more modest Monte-Carlo approach that allowed us to infer the probability of operational convergence.

In the case of a 4-segment aperture with piston, x-tilt and y-tilt misalignments there are a total of nine independent parameters to estimate. A representative sampling of the parameter space is required for a sound Monte Carlo experiment. The nine-dimensional parameter space was sampled based on the value of the RMS phase error of the misalignment. This approach has some advantages. First, an alternative, uniform sampling of the nine-dimensional parameter space would require the testing

of tens of thousands of misalignment parameter sets, requiring months of computing. By approaching this problem with respect to the RMS error of the random misalignments, we can significantly reduce the number of simulations.

Second, because misalignments with an RMS phase error below the Marechal aberration-tolerance condition of $1/14^{\text{th}}$ of a wave (0.0714 wave) [3.2] already produce images of sufficient quality, we can avoid testing the global convergence of misalignments satisfying this condition. In addition, the condition provides a tolerance for "operational" convergence. The Marechal condition insures that the intensity at the diffraction focus is at least 0.8 (= Strehl intensity) times that given by a diffraction-limited system. This relationship between the Marechal condition and the Strehl intensity is independent of the design of the binary pupil function and therefore applies to phased-array systems. Systems that meet the Marechal condition are considered to be well-corrected, giving imagery for which the degradation would be difficult to perceive.

A third advantage to sampling the parameter space with respect to RMS phase error is that the severity of image degradation typically increases as this error increases. Thus the operational convergence is naturally linked to image fidelity via RMS misalignment error. By testing algorithm convergence with respect to RMS phase error, we can gain an understanding for just how much misalignment can be tolerated before the probability of operational convergence is no longer acceptable. This provides a measure of the phase-diversity capture range, or misalignment-acquisition ability. The issue of accurately tracking the MMTT alignment in an adaptive system also is addressed by this approach. The probability of operational convergence allows us to predict how far out of alignment the telescope can drift before phase-diversity can be reliably used to realign it. Furthermore, if the parameter sets to be tested are carefully chosen over a wide range of RMS phase errors, then through interpolation we can estimate the probability of operational convergence for misalignments that are not directly tested.

3.2 EXPERIMENT DESIGN

The first step in designing the probability-of-convergence experiment was to estimate the probability density of the RMS phase error for random MMTT misalignments.

To do this we generated random MMTT misalignments according to the following probability densities: (1) piston errors are uniformly distributed from $-\lambda/2$ and $\lambda/2$, and (2) x- and y-direction tilts are uniformly distributed from $-\lambda/3$ and $\lambda/3$, where λ represents the wavelength of the illumination. We have assumed narrow-band illumination throughout, which is a valid assumption for intended laboratory demonstration experiments. These aberration intervals are based on the expected range-of-motion of the mirror segments as prescribed by Phillips Laboratory engineers. Thus piston errors are permitted over a full, single-wave range, and tilt errors are constrained over a more limited range. When generating random misalignments, all of the mirror segments were free to move. As far as the phase-diversity algorithm is concerned, there are only nine free parameters, and one of the mirror segments can be considered fixed. However, we decided it would be more realistic to investigate the misalignments that typically arise when all the segments are free to move.

It is difficult to estimate the probability density function (PDF) for RMS misalignment error analytically, due to global tilt and piston ambiguity issues [3.1]. Therefore, the RMS phase error was calculated for 5-million random sets of misalignments, and the resulting estimate of the PDF for RMS misalignment error is shown in Figure 3-1. As is appropriate for a PDF, this plot has been normalized to unit area. The mean RMS misalignment is approximately 0.16 waves, which is quite close to the location of the peak of the density. No random misalignments were found with an RMS error less than 0.02 or greater than 0.27 waves. The probability of a random misalignment having an RMS error less than the Marechal criterion of $\lambda/14$ is about 0.1%; thus, given uniformly-distributed misalignment parameters, it is highly likely that real-time or post-detection aberration correction will be needed. It should be reiterated that the PDF in Figure 3-1 is based on the assumption of uniformly-distributed random piston and tilt parameters and should not be considered a refined model for expected misalignments once the telescope is operational. The total probability of operational convergence (a single number) is found by multiplying the RMS-misalignment PDF by the probability of operational convergence *as a function of RMS misalignment error* and then integrating over RMS error. It remains to choose the number of RMS bins and the number of parameter sets to test within each bin.

The RMS-misalignment PDF aids us in the design of the Monte-Carlo experiment. The ultimate goal is an understanding of the phase-diversity probability of operational




Figure 3-1: Estimate of PDF for MMTT RMS Misalignment Error. We assume piston errors uniformly distributed on $[-\lambda/2, \lambda/2]$ and x- and y-tilt errors uniformly distributed on $[-\lambda/3, \lambda/3]$. The six bins used in the probability-of-convergence experiment are shown.

convergence. Choosing regions of interest over the range of the density function enables us to investigate the convergence rate in terms of the RMS phase error of the misalignment. We selected six RMS bins with a very small width of 0.0025 waves and generated random misalignment parameters with RMS phase errors that fall within these bins. These bins are located at RMS misalignment levels of 0.075, 0.0875, 0.10, 0.15, 0.20, and 0.23, as shown in Figure 3-1. Given a particular bin and a large number of misalignment parameter sets that fall into this bin, the next step is to apply phase diversity to each of the parameter sets. The percentage of parameter sets for which the algorithm converges to an acceptable solution provides an estimate of the probability of operational convergence for misalignments within the particular range of RMS error defined by the bin location and size. By obtaining estimates of the probability of convergence for a handful of bins spanning the range of the PDF, we can interpolate to roughly predict the probability of operational convergence for all points within this range. This effectively reduces the number of trials of the Monte-Carlo experiment.

3.3 EXPERIMENTAL ASSUMPTIONS

Perhaps the most important experimental assumption is the following: *all phase-diversity parameter estimates had a starting guess of zero misalignment.* This is quite important for the following reason. Convergence to the global maximum is highly dependent on where the search is started. Thus, any measure of the probability of finding the global maximum is also highly dependent on the starting guess. We designed the experiment with the worst-case scenario in mind: the telescope is misaligned and there is no *a priori* knowledge available about how badly it is misaligned, which mirror segment(s) has moved, what the state of the telescope was after the last alignment, etc. Beginning the search with a zero misalignment implies an assumption that the telescope is already aligned. Because of this, the result of the experiment will provide information about the capture range of the algorithm; i.e., how much misalignment is tolerable before the probability of operational convergence falls below some desired bound. It also gives an indication of how much misalignment is tolerable for accurate real-time tracking.

A definition of "operational convergence" is also necessary. A maximization se-

quence is defined to achieve convergence when the RMS phase error of the estimate with respect to the true misalignment falls below the Marechal condition of $1/14^{th}$ of a wave. Stagnation is verified by restarting the gradient-search algorithm twice after the stopping criterion has been initially satisfied. The stopping criterion is based on a percentage change in the objective function, and the effect of restarting the search is to reset the search vector in the direction of steepest ascent. If the search again stagnates, and the RMS error of the estimate is above $1/14^{th}$ of a wave, then we say that convergence to an acceptable estimate has not been achieved.

3.4 REVIEW OF NOISELESS EXPERIMENT

A 128 x 128 image of a jet was used as the object for the noiseless operational convergence experiments. Therefore, the estimated probability of operational convergence technically applies only to this single image. Given additional resources, the Monte-Carlo experiments could easily be expanded to include a variety of scenes. Still, the single-image results provide an important flashlight into more general operational convergence.

Two wavelengths of quadratic diversity were used, and the MMTT fixed aberration was not added because this would have no effect on misalignment estimates if included in the imaging model [3.1]. A total of 600 parameter sets for each of the six RMS phase bins were tested for convergence. For a given range of RMS phase, each experiment has a binary outcome (operational convergence or stagnation), so that the number of successful convergences for each RMS bin is analogous to the number of times a weighted coin turns up heads in a series of coin-toss experiments. Therefore, the number of successful convergences will follow a binomial probability law

$$P_N(m) = \frac{N!}{(N-m)!m!} p^m (1-p)^{N-m}, \quad (3-1)$$

where p is the probability of operational convergence, N is the number of trials, and m is the number of successful convergences. We are interested in determining the parameter p , which has a maximum-likelihood estimate

$$\hat{p} = m/N, \quad (3-2)$$

which is just the percentage of operational convergences in N trials. The standard deviation of the estimate \hat{p} is

$$\sigma_{\hat{p}} = (p(1 - p)/N)^{1/2}, \quad (3 - 3)$$

which reflects the uncertainty of the estimate of the probability of convergence. For example, the standard deviation of \hat{p} for 600 trials of a binomial experiment with probability $p = 0.9$ is 0.01. On the other hand, if $p = 0.5$, then $\sigma_{\hat{p}}$ is a maximum with a value $\sigma_{\hat{p}} = 0.02$. Therefore 600 trials per RMS bin will provide an estimate of the probability of operational convergence in which we can have high confidence.

The results from the experiment are shown in Figure 3-2, which also has the RMS PDF superimposed for reference. We expected the convergence rate to drop as the amount of misalignment moved farther away from our initial guess of zero misalignment. The worst case was the $[0.23, 0.2325]$ bin at the edge of the distribution, yet there is still an 89.5% success rate for these most severe aberrations. We can conclude from this experiment that, given the assumption of uniformly-distributed piston and tilt misalignments over the ranges specified above and an initial guess at the origin of parameter space, the probability of operational convergence for the most-expected types of misalignments is above 90%.

We conjecture that stagnation is primarily due to objective-function surfaces for which the initial estimate of an aligned system (naive initial estimate) falls within the capture range of a local maximum. When a naive initial estimate stagnates, it is quite possible that another initial estimate will avoid the local maximum and yield an acceptable convergence. A polling approach that selects the misalignment estimate from estimates computed in parallel and obtained from various initial guesses could be used. The hope is that the probability of operational convergence for a given misalignment and as a function initial estimate will be roughly comparable to that for a fixed initial estimate and as a function of misalignment realization. If this turns out to be true, then only a few initial estimates used in parallel could increase the operational convergence dramatically.

A rate of 100% convergence was obtained for the 0.085 bin (close to the 0.0714 Marechel criterion), and even at 0.1 wave of RMS error a global convergence rate of 98.8% was observed. An RMS phase capture range of the algorithm can be estimated




Figure 3-2: MMTT Probability of Operational Convergence with No Noise. The upper curve is the probability of operational convergence for a given RMS bin, and the lower curve is the PDF of Figure 3-1. The overall probability of operational convergence would be obtained by multiplying these two curves and integrating over all possible RMS values. Piston-plus-tilt misalignment errors were used.

via interpolation for a desired probability of operational convergence. If applied in a real-time tracking mode, the phase-diversity algorithm could be used to update the alignment before the RMS error strays too far beyond the "knee" in the curve, perhaps around $\text{RMS} = 0.1$ wave. The random misalignment model here is an unlikely scenario for alignment degradation over short periods of time. From this model the probability of obtaining a random misalignment with an RMS less than 0.1 wave is only about 2%. This probability would likely be much larger for a system that "gracefully" drifts out of alignment. The probability-of-convergence curve is useful in this case for estimating the capture range as the telescope gradually misaligns. In conclusion, the likelihood of algorithm stagnation is quite small even for the worst-case scenario of no *a priori* knowledge of misalignment, and for the ideal case of noiseless imaging we expect the algorithm to find an operationally acceptable solution with high probability.

3.5 OPERATIONAL CONVERGENCE WITH NOISE

The results of the noiseless phase-diversity convergence experiments are quite useful, but not necessarily applicable to real MMTT operation. Certainly if we found the probability of operational convergence to be unacceptable in the noiseless case, then we would only expect performance to degrade when CCD noise was added. Because the noise-free performance was so good, it was worth pursuing additional experiments that included noise. We conducted the same experiment in the presence of noise with the hope of adding more realism to our estimates of the probability of operational convergence for the MMTT.

We were unable to obtain specific information about the expected signal-to-noise ratio in the MMTT CCD arrays, so we used some general knowledge about CCDs in arriving at the appropriate noise level. The signal-to-noise ratio in a CCD will be limited by the number of sampled photoelectrons in each detector element, assuming that readout noise can be neglected. CCD cameras have a full-well depth of about 4×10^5 photoelectrons. At room temperature, a medium-sized format CCD array ($\simeq 400$ pixels on a side) can operate at approximately 60 Hz with photon sampling being the dominant noise, so long as the number of photoelectrons is a reasonable percentage of the full well [3.3]. Assuming that a given detector element records

half full well, the signal-to-noise ratio will be 4.47×10^2 . Stated differently, the noise will be about 0.2% of the signal. For a low-contrast object, this noise is well approximated with a Gaussian noise model. Improved signal-to-noise ratios can be achieved by averaging frames, so long as the imagery does not change at the frame rates employed.

We feel that the signal-to-noise figure specified for the CCD arrays is fairly high and that a more pessimistic figure would be appropriate for these experiments. Rather than assuming a signal-to-noise of about 450, we cut this figure roughly in half to about 200. In other words, we assumed the expected noise to be about 0.5% of the signal. Also, the standard deviation of the noise for all the detector elements was set to 0.5% of the *peak* value in the diffraction-limited jet image. Thus, the Gaussian approximation to photoelectron noise is responsible for a larger-than-expected noise level at those detector elements receiving a small number of photons. For a good CCD array and a strong signal, this noise level is therefore pessimistic. However, if the phase-diversity probability of operational convergence is acceptable for this higher noise level, then we can expect even better performance for CCD arrays with better signal-to-noise ratios. These noise levels are appropriate for laboratory demonstrations, however photon levels that would be expected in an actual field application have not yet been determined.

3.6 EXPERIMENTAL RESULTS

To facilitate comparisons with the noiseless probability-of-convergence experiment, the noisy experiment was performed using the same 6 bins and the same 600 misalignment parameter sets within each bin. We have observed in previous work [3.1] that the accuracy of the parameter estimation can vary not only with noise level, but also with the realization of the noise at a given level. A unique noise realization was added to every simulated image in the experiment, adding another dimension of randomness to the experiment. The regularization parameter was selected as the ratio of the variance of the noise to the variance of the collected conventional image. This results in a slightly larger parameter than would be obtained when using the variance of the object. However, our experience has been that estimates are relatively insensitive to comparable variation in the regularization parameter.

The portability of the software and the parallel nature of the Monte-Carlo computations enabled us to utilize the computing power of an entire network of resource-sharing workstations. A simple extension of the C language enabled us to use nearly 20 Sun workstations of various speeds as a distributed parallel processor. Each processor works on an assigned parameter set, sends its results to a central location, and begins work on a new parameter set until none remain. This has the advantage of bringing to bear machines of various speeds without a significant loss in throughput due to one or more particularly slow processors. For this task we processed the 3600 parameter sets nearly 18 times faster than if a single workstation had been tasked.

The results from two of the bins are shown in Figure 3-3. This figure plots the probability of operational convergence versus the convergence threshold for two bins: the 0.075 and 0.23 RMS phase error bins. These plots contain all the information from the experiment for a particular bin. For our definition of operational convergence, we have chosen the Marechal condition of 0.0714 wave RMS error. Figure 3-3 indicates that, for this condition, the probability of operational convergence for the 0.075 bin is 1, whereas the probability for the 0.23 bin is about 0.66 (roughly two out of every three). If the operational convergence criterion were made more stringent and set to 0.05 wave, then the probability of operational convergence for these two bins would drop to 84% and 42%, respectively.

With the convergence threshold set to the Marechal condition, we added the probability of operational convergence for the noisy experiment to Figure 3-2. This plot, shown in Figure 3-4, demonstrates the performance loss associated with the addition of noise to the collected images. The probability of convergence is lower for the noisy experiment in all but the 0.075 bin. The probability drops to 87% in the 0.10 bin and then drops dramatically to 44% for the 0.15 bin. We were surprised that the probability then increases for the two largest bins at 0.2 and 0.23. The approximately monotonic behavior observed in the noiseless curve is not demonstrated in the noisy curve as we might have expected. We would need additional samples on this curve to better understand this behavior, and we currently have no explanation for this trend.

The overall probability of convergence can be estimated by multiplying the probability of convergence curve with the RMS phase error probability density and integrating. The overall probability is roughly 50%. This probability is still favorable for MMTT startup. That is, the phase-diversity algorithm will estimate the misalign-




Figure 3-3: Probability of Operational Convergence Versus Convergence Threshold. Two RMS bins are considered and the simulated MMTT data included noise.

ment to within the Marechal condition half the time when starting up the system. In terms of tracking the misalignment during operation, the probability of convergence curve offers insight into how far out of alignment the MMTT should be allowed to drift. The probability of convergence is greater than 95% when the misalignment is less than 0.08 wave RMS, just above the Marechal condition. A more relaxed tracking could allow the system to drift to between 0.1 and 0.12 wave RMS misalignment with a greater than 70% probability of realignment via phase diversity. In cases where the misalignment estimate error is greater than the Marechal condition, the system could be realigned to these parameters, reimaging could be performed, and a new set of misalignment parameters could be estimated with a higher probability of convergence. Alternatively, multiple initial estimates could be used in parallel to improve the probability of operational convergence. Strategies for selecting the number and values of multiple initial estimates and quantifying the corresponding improvement in probability of operational convergence are important issues for future research.




Figure 3-4: MMTT Probability of Operational Convergence. The upper curve is the probability of operational convergence for a given RMS bin with no noise, and the lower curve is the probability of convergence for a given RMS bin with noise. The probability density function of Figure 3-1 is overlaid for reference. The overall probability of operational convergence would be obtained by multiplying a probability of operational convergence curve with the probability density and integrating over all possible RMS values.

3.7 MONTE-CARLO EXPERIMENT SUMMARY

A Monte-Carlo approach allowed us to infer the probability of operational convergence for the case of piston-plus-tilt misalignments and noisy collected images. Quantifying misalignments by RMS phase error enabled us to conveniently characterize the phase-diversity capture range and tracking ability. By assuming a smooth degradation with increasing RMS phase error, we were able to infer the probability of convergence for all possible parameters from experiments on a small subset of these parameters. We compared the results of this experiment with an identical experiment with no noise and found that overall probability of convergence is less for the noisy case. The probability of convergence did not exhibit the monotonic decreasing behavior seen in the noiseless experiment, and the probability in the 0.15 RMS error bin dropped by more than one-half. This behavior is not fully understood, and we feel that more samples

along this curve would provide additional insight. The overall probability of convergence was estimated at between 50 and 55% as compared with an approximate 94% convergence rate for the noiseless case. Tracking the misalignment with greater than 70% probability can be achieved by keeping the MMTT from drifting beyond 0.12 wave RMS misalignment. Overall we have shown that the likelihood of operational convergence for the worst-case scenario of no *a priori* knowledge of misalignment is smaller in the presence of noise, but we feel that the startup and tracking probability is sufficiently high in the regime of bright (photon-rich) scenes, particularly in light of the fact that the noise levels were set to slightly pessimistic levels given this regime.

In Section 4.3 we show that an additional diversity image can be used to dramatically increase the probability of operational convergence in the case of noisy data.

REFERENCES

- [3.1] R.G. Paxman, J.H. Seldin, and P.P. Sanchez, "Applied Phase Diversity," final technical report to Phillips Laboratory (Kirtland Air Force Base), ERIM Report No. 213390-3-F, March 1992.
- [3.2] M. Born and E. Wolf, *Principles of Optics* (Pergamon Press, New York, 1980) pp. 468,469.
- [3.3] B.W. Lites, "Sensitivity to polarization: How can we measure weak vector magnetic fields at high angular resolution?", in *Solar Polarimetry*, Proceedings of the 11th Sacramento Peak Summer Workshop, L.J. November ed., Sunspot, New Mexico, p. 173 (1991).

4.0 MULTIPLE DIVERSITY FRAMES

The generalization of the Gonsalvez two-image phase-diversity method [4.1,4.2] to multiple diversity images is presented in a recent paper [4.3]. We have traditionally used just two diversity images in MMTT phase-diversity research, and we discuss here the use of additional diversity images for improving MMTT-based misalignment and object estimation. An example of this improvement in severe noise conditions is shown, and the Monte-Carlo noise experiment discussed in Section 3 is modified to use 3 diversity images, resulting in a dramatic improvement in the probability of operational convergence.

4.1 POTENTIAL ADVANTAGES OF MULTIPLE DIVERSITY FRAMES

The collection of more than two diversity images (frames) offers several potential advantages. With the addition of diversity images the number of quantities to estimate remains unchanged (the object and a set of aberration parameters), but the amount of information available to perform the estimation is significantly increased. More observations of quantities related in a known way to the desired estimates will usually result in lower estimate errors. This is of particular importance when collecting data that are corrupted by noise. Additional diversity images can be thought of as collecting new information from another "look" or "perspective," reducing ambiguities created by noise or by insufficient diversity. Another advantage is that the collection of a series of images in different focal planes provides improved sampling of the focal volume, which is a viewpoint considered by Högbom [4.4]. Better estimates of aberration parameters could significantly improve the probability of operational convergence. An experiment to verify this supposition is discussed in the next section.

With improved aberration estimates also comes the ability to improve the object estimates. Not only will the object estimate improve due to the fact the misalignments are known more precisely, but the restoration will be aided by additional MTF fill-in. Additional diversity can supplement spatial frequency information that otherwise may have been missing or severely attenuated. These additional spatial frequencies, which are typically high frequencies, produce sharper object estimates with more information. In the case of a sufficiently bright (photon-rich) object for which the

detector dynamic range limits the signal-to-noise on all detector arrays, the signal collected with each additional frame of diversity also aids the suppression of noise in the estimated object.

4.2 THREE-FRAME DIVERSITY EXPERIMENT

The addition of diversity frames in an operational system could be hindered by a number of factors, including cost, power, space, and available light. With this in mind we decided to explore the simplest case of a single additional diversity image. We first identified a set of misalignment parameters which, in the presence of noise, had a 2-frame objective-function with an apparent global maximum that predicted a very poor misalignment estimate. A misalignment estimate is presumed to be a global maximum of the objective function if it is identical to the misalignment estimated from a search that begins at the true solution. In the case of no noise, a conjugate-gradient search that begins at the true solution would stagnate immediately at the true solution; however, when noise is present in the imagery, the search often moves away from the true solution to a set of misalignment parameters that give rise to a smaller value of the objective function.

The simulated images from this set of misalignment parameters are shown in Figure 4-1. Using the ratio of the image variance to the noise variance, the signal-to-noise ratio of the collected images is 20 dB. Stated another way, the ratio of the standard deviation of the noise to the image peak is 2.24%, roughly five times more noise than used for the Monte-Carlo probability of operational convergence experiment. This is a much larger noise level than expected for a bright object and a good CCD array, but it presents the phase-diversity algorithm with a greater challenge and a larger potential gain. Figures 4-1(A) and (B) present the conventional and 2-wave quadratic diversity images, respectively, simulating the MMTT with a 0.2 wave RMS misalignment. The objective function defined by these images and the MMTT system has an apparent global maximum that yields a misalignment estimate with a 0.11 wave RMS error with respect to the true solution. The object estimate for this 2-frame misalignment estimate, the Fourier transform of which is defined in Eq. (2-18) for $K = 2$, is shown in Fig. 4-1(D). Residual misalignment phase errors are manifest in the blurred edges along the wings, in the camouflage areas, and at the shadow


boundaries.

Another maximization from the true solution to find the global maximum was performed with the same misalignment, but with an additional diversity image. This image with -1 wave of quadratic diversity is shown in Figure 4-1(C). With the objective function utilizing all three frames in the first row of Figure 4-1, a misalignment estimate with a residual error of only 0.068 wave was found. The 3-frame maximum-likelihood restoration of the object is shown in Figure 4-1(E). This image is far superior to the 2-frame object estimate in part (D) for two main reasons. The first reason is that, unlike the 2-frame estimate, the misalignment estimate error is below the Marechal condition, and an error of this size does not have significant visual impact on the image quality. Second, an additional diversity frame provides an overall sharper estimate due to additional MTF fill-in and some averaging out of noise at all spatial frequencies. In fact, we find that the 3-frame restoration with a 0.068 wave residual RMS phase error is superior yet to the ideal 2-frame restoration from an aligned system displayed in Figure 4-1(F). A comparison of Figures 4-1(E) and (F) clearly shows that, for a strong signal source, additional diversity improves the estimate of the object. We feel that this improvement will be significant regardless of the level of noise in the collected images.


4.3 THREE-FRAME PROBABILITY OF OPERATIONAL CONVERGENCE

The noisy Monte-Carlo probability of convergence experiments discussed in Section 3.5 revealed a substantial reduction in performance over the case of no noise. In particular, the three largest RMS phase error bins (0.15, 0.2, and 0.23) exhibited the largest drop in probability of convergence with the addition of noise. We performed the same experiment with 500 of the 600 parameters in these three bins, but with an additional frame with -1 wave quadratic diversity. Our goal was to investigate the impact of a single additional diversity image on operational convergence.


The results of the 3-frame Monte Carlo experiment, shown in Figure 4-2, were impressive. The 3-frame probability of operational convergence is the dashed curve in Figure 4-2. For these 3 bins, the probability of operational convergence using the Marechal criterion improved in all cases to greater than 90%. The 3-frame probability




(A)



(B)




(C)



(D)



(E)



(F)

Figure 4-1: Three-Frame Phase Diversity. (A) Noisy conventional image with 0.2 wave RMS misalignment; (B) Noisy diversity image 1 with 2 waves quadratic diversity; (C) Noisy diversity image 2 with -1 wave quadratic diversity; (D) 2-frame restoration using misalignment estimate with 0.11 wave RMS error; (E) 3-frame restoration using misalignment estimate with 0.068 wave RMS error; (F) Ideal 2-frame restoration from perfectly aligned system.




Figure 4-2: MMTT Two- and Three-Frame Probability of Operational Convergence. The upper solid curve is the 2-frame probability of operational convergence for a given RMS bin with no noise, and the lower solid curve is the 2-frame probability of operational convergence for a given RMS bin with noise. The 3-frame with noise curve is shown for the largest three RMS bins. The PDF of Figure 3-1 is overlaid for reference.

of operational convergence for the 0.15 RMS phase bin is roughly double that of the 2-frame counterpart, the 0.2 RMS bin approximately matches the 2-frame noiseless probability, and at 96% the probability of operational convergence for the 0.23 RMS phase bin exceeds the corresponding 2-frame, no noise probability. The implication is that a single additional diversity image will significantly improve MMTT startup from an unknown misalignment, and that realignment update intervals in the real-time tracking mode could be relaxed. From these results combined with the results from the previous section, we conclude that object and misalignment parameter estimates will be enhanced by collecting a single additional diversity frame. The value added by collecting frames for $K > 3$ frames has not been explored.

4.4 SUMMARY OF MULTIPLE-FRAME PHASE DIVERSITY

We have demonstrated the ability of the phase-diversity software to simulate multiple-frame phase diversity, and we have recounted two experiments using 3 diversity images. The first experiment was an example of how, under very noisy conditions, an additional diversity image was capable of improving the misalignment estimate so as to satisfy the Marechal criterion. This experiment also demonstrated the improvement in quality of the object estimate obtained with an additional frame. The object estimate from 3 frames of data derived from a misaligned system was still better than the restoration from 2 frames of data derived from an aligned system. The second experiment, an expansion of the 2-frame noisy Monte-Carlo probability of convergence experiment discussed in Section 3.5, was an exhibition of the improvement in misalignment estimate accuracy provided by a single additional diversity image. The 3-frame probability of operational convergence was significantly improved to greater than 90% for misalignments with a 66% or less convergence rate in the 2-frame case. This dramatic increase has implications for improved MMTT start-up accuracy and more relaxed real-time tracking.

REFERENCES

- [4.1] R.A. Gonsalvez and R. Childlaw, "Wavefront sensing by phase retrieval," in *Applications of Digital Image Processing III*, A.G. Tescher, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **207**, 32-39 (1979).
- [4.2] R.A. Gonsalvez, "Phase retrieval and diversity in adaptive optics," *Opt. Eng.* **21**, 829-832 (1982).
- [4.3] R.G. Paxman, T.J. Schulz, and J.R. Fienup, "Joint estimation of object and aberrations by using phase diversity," *J. Opt. Soc. Am. A* **9**, 1072-1085 (1992).
- [4.4] J.A. Högbom, "On the intensity distribution over the focal volume," in *High Spatial Resolution Solar Observation*, O. von der Lühe, ed., proceedings of the Tenth Sacramento Peak Summer Workshop, Sunspot, NM (August, 1988).

APPENDIX A

APPLIED PHASE DIVERSITY SOFTWARE

UPGRADES

This appendix gives an overview of the current state of the UNIX workstation-based Applied Phase Diversity (APD) software. Upgrades to the software that improved both the speed and accuracy of phase-diversity misalignment estimation are discussed along with some new features and capabilities.

A.1 APD OVERVIEW

The phase-diversity algorithm is currently implemented in a piece of software called APD. The program is basically interactive, with a command-line lexicographical interpreter delegating the requested tasks. The program can also be run in a batch mode and has a session-logging capability that is useful for debugging and experiment documentation. The software is currently tailored to a segmented-aperture telescope with piston and tilt misalignments, yet it is written in a modular fashion that enables an easy adaptation to incoherent imaging systems with monolithic apertures and/or different misalignment parameterization schemes. The program can perform simulations of an imaging system using ideal, digital objects and known misalignments, or it can emulate the imaging process and use real, collected data to estimate the unknown misalignment parameters. The APD software was used primarily in the first mode to perform simulations of the MMTT. This enabled us to test the misalignment-estimation performance of the phase-diversity algorithm for various misalignment parameters, amounts and types of diversity, detector noise levels, objective-function regularization parameters, etc. A more detailed discussion of the basic capabilities of the APD package can be found in Appendix C in [A.1].

A.2 APD UPGRADES

In this section we present the upgrades to the APD software that were accomplished in the current effort.

A.2.1 Analytic Gradient

Misalignment estimation is performed by maximizing the phase-diversity objective function. In practice we minimize the negative of this objective function using a conjugate-gradient algorithm. Past versions of the APD utility implemented a two-sided finite-difference approximation to the gradient of the objective function. An analytic gradient was derived in [A.1] and was implemented in software later as a replacement for the finite-difference approximation. Not only is the analytic form a more accurate means for calculating the gradient, but it also provides a great computational savings. In the case of a J-parameter minimization using K diversity images, the two-sided finite-difference method requires $4JK$ FFTs for each evaluation of the gradient. The analytic gradient reduces the number of FFTs to $4K$. For the case of two diversity images and 9 independent parameters (piston, x- and y-tilt for a 4-segment aperture), the finite-difference estimate would require 72 FFTs while the analytic gradient would compute only 8 FFTs. The evaluation of the analytic gradient requires more non-FFT computations than does the finite-difference method; however, the FFT is the most computationally intensive portion of the calculation, and the gain in using the analytic gradient is still considerable. Note as well that the number of FFTs required by the analytic gradient is independent of the number of parameters being estimated. Because of its direct dependence on the number of parameters, the finite-difference technique would not be a feasible approach for large J. As an example, the use of finite differences for a point-by-point misalignment characterization would be computationally prohibitive.

A.2.2 Multiple Diversity Images

The previous version of the software assumed that there would be a single diversity image to complement the conventional focal-plane image. However, the phase-diversity theory generalizes to include an arbitrary number, K, of collected images; thus, the APD software was upgraded to facilitate K-image, or multi-frame, phase diversity. Each diversity image can have any combination of linear (piston and/or tilt) or quadratic diversity. The generalization to K images manifests itself in all parts of the phase-diversity software, including the objective function evaluation, the analytic gradient calculation, and the object estimation (see below). Multiple-frame

diversity experiments are discussed in Section 4, where we conclude that even a single additional diversity image can have a significant impact on both parameter estimation accuracy and object estimate quality and can also help to improve the capture range of the phase-diversity algorithm.

A.2.3 Object Estimation

Given a fixed aberration function, a closed-form expression exists for the underlying object that maximizes the regularized log-likelihood objective function. This expression of the regularized maximum-likelihood object estimate is the natural means for obtaining a post-detection restoration of the image data. The APD software allows the user to save the object estimate for any set of misalignment estimates and any regularization parameter. This feature complements the other APD image-save option that assumes a real-time realignment of the optics followed by a reimaging of the object.

A.2.4 Numerical Improvements

Two significant numerical improvements were made to the APD code. The first change enables the user to choose between a single-precision and a double-precision implementation of the software. We consistently used the double-precision version for our simulations. The single-precision version runs with less allocated memory and may be preferable in certain situations. No study of the effect of precision on misalignment estimation has been performed.

The second improvement was the result of an investigation of round-off errors when summing a set of numbers spanning a large dynamic range. The summation under consideration is the sum over the spatial-frequency variable u in the objective function. The frequency-domain quantities being summed span a very large dynamic range, with a handful of very large numbers at the low spatial-frequencies (DC does not change over the misalignment parameters and thus can be ignored). Severe round-off errors are likely if these very few large terms are included in the summation prior to adding the smaller terms, which account for the majority of the numbers. To minimize round-off error, one ideally would sort the summands in ascending order and then would sum over u starting with the smallest value.

This sorting procedure is quite expensive computationally, so a sub-optimal alternative that exploits the low-pass nature of an incoherent imaging system was employed. Rather than sorting the summands, the sum is performed in the spatial-frequency domain roughly along semi-circular arcs about the DC point. We start along the arc farthest from DC and gradually move inward toward the large values at the low spatial frequencies. We investigated this technique and found that the resulting residual round-off errors were insignificant and that the results compared favorably with a pre-sorted summation.

A.3 SUMMARY

The latest version of the phase-diversity algorithm has some new benefits and makes some improvements on the previous version. The analytic gradient provides a faster and more accurate calculation of a quantity that is crucial to parameter estimation via conjugate-gradient objective-function minimization. The ability to simulate the collection of multiple (greater than two) diversity images is important for analyzing the gain in misalignment estimate accuracy, the effect on phase-diversity capture range, and the improvement in object estimation. Improvements in numerical accuracy make the phase diversity software a more reliable tool.

REFERENCES

- [A.1] R.G. Paxman, J.H. Seldin, and P.P. Sanchez, "Applied Phase Diversity," final technical report to Phillips Laboratory (Kirtland Air Force Base), ERIM Report No. 213390-3-F, March 1992.